shifted upward, therefore no spectral overlapping occurs. If it is an NTSC channel, the shift is below the (RF equivalent of the) Nyquist slope of an NTSC receiver where there is high attenuation, and it is slightly above its customary lower adjacent-channel sound trap. No adverse effects of the shift have been found, nor are they foreseen. An additional shift of the ATV spectrum is used in order to track the dominant NTSC interferer which may be assigned an offset of -10 kHz, 0 kHz or + 10 kHz.

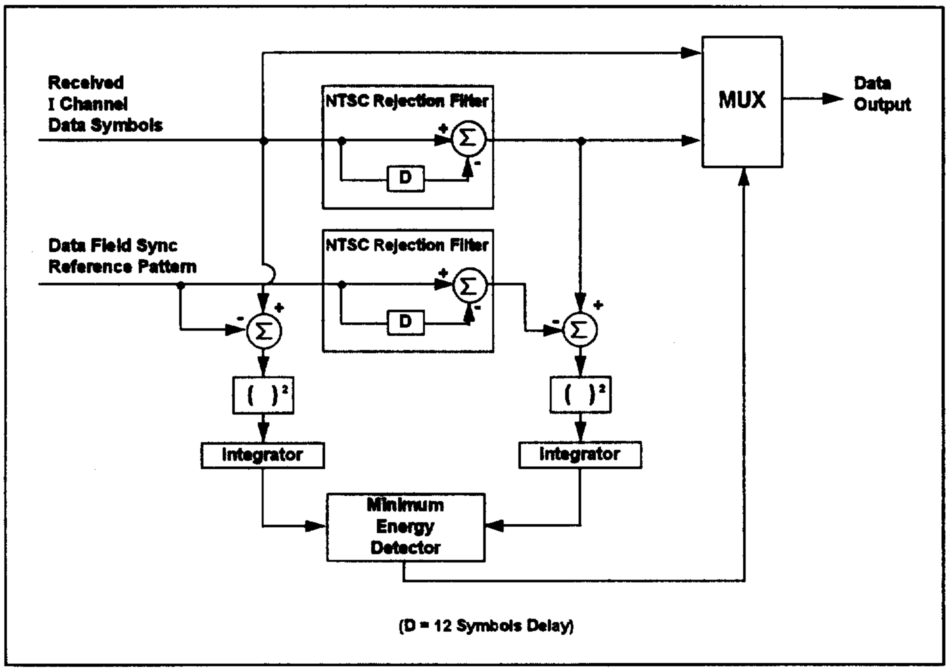

Figure 10.8. NTSC interference rejection filter.

NTSC interference can be detected by the circuit shown in Figure 10.8, where the signal-to-interference plus noise ratio of the binary Data Field Sync is measured at the input and output of the comb filter, and compared to each other. This is accomplished by creating two error signals. The first is created by comparing the received signal with a stored reference of the field sync. The second is created by comparing the rejection filter output with a combed version of the internally stored reference field sync. The errors are squared and integrated. After a predetermined level of confidence is achieved, the path with the largest signal-to-noise ratio (lowest interference energy) is switched in and out of the system automatically.

There is a reason to not leave the rejection comb filter switched in all the time. The comb filter, while providing needed co-channel interference benefits, degrades white noise performance by 3 dB. This is due to the fact that the filter output is the subtraction of two full gain paths, and as white noise is un-correlated from symbol to symbol, the noise power doubles. There is an additional 0.3 dB degradation due to the 12 symbol differential coding. (See Section 10.2.3.9, Trellis decoder.) If little or no NTSC interference is present, the comb filter is automatically switched out of the data path. When the NTSC service is phased out, the comb filter can be omitted from digital television receivers.

10.2.3.7 Channel equalizer

The equalizer/ghost canceller compensates for linear channel distortions, such as tilt and ghosts. These distortions can come from the transmission channel or from imperfect components within the receiver.

The equalizer uses a Least-Mean-Square (LMS) algorithm and can adapt on the transmitted binary Training Sequence as well as on the random data. The LMS algorithm computes how to adjust the filter taps in order to reduce the error present at the output of the equalizer. It does this by generating an estimate of the error present in the output signal. This error signal is used to compute a cross-correlation with various delayed data signals. These correlations correspond to the adjustment that needs to be made for each tap to reduce the error at the output.

The equalizer algorithm can achieve equalization through three means: it can adapt on the binary training sequence; it can adapt on data symbols throughout the frame when the eyes are open; or, it can adapt on data when the eyes are closed (blind equalization). The principal difference among these three methods is how the error estimate is generated. (See Section 10.3, Receiver equalization issues.)

For adapting on the training sequence, the training signal presents a fixed data pattern in the data stream. Because the data pattern is known, the exact error is generated by subtracting the training sequence from the output.

The training sequence alone, however, may not be enough to track dynamic ghosts as these require tap adjustments more often than the training sequence is transmitted. Therefore, once equalization is achieved, the equalizer can switch to adapting on data symbols throughout the frame, and produce an accurate error estimate by slicing the data with an 8-level slicer and subtracting it from the output signal.

For fast dynamic ghosts (e.g., airplane flutter) it is necessary to use a blind equalization mode to aid in acquisition of the signal. Blind equalization models the multi-level signal as binary data signal plus noise, and the equalizer produces the error estimate by detecting the sign of the output signal and subtracting a (scaled) binary signal from the output to generate the error estimate.

To perform the LMS algorithm, the error estimate (produced using the training sequence, 8-level slicer, or the binary slicer) is multiplied by delayed copies of the signal. The delay depends upon which tap of the filter is being updated. This multiplication produces a cross-correction between the error signal and the data signal. The size of the correlation corresponds to the amplitude of the residual ghost present at the output of the equalizer and indicates how to adjust the tap to reduce the error at the output.

A block diagram of the equalizer is shown in Figure 10.9. The DC bias of the input signal is first removed by subtraction. The DC may be caused by circuit offsets, non-linearities, or shifts in the pilot caused by ghosts. The DC offset is tracked by measuring the DC value of the training signal.

The equalizer filter consists of two parts, a 64 tap feed-forward transversal filter followed by a 192 tap decision feedback filter. The equalizer operates at the 10.762 MHz symbol rate (T-sampled equalizer).

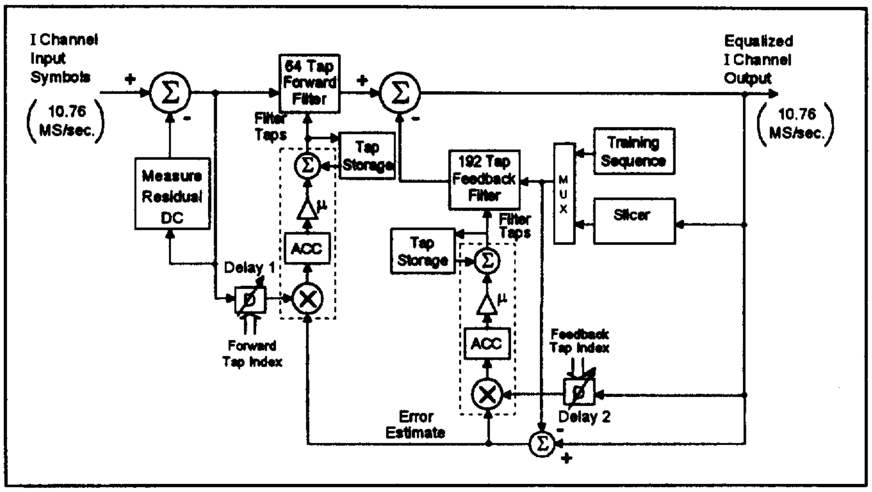

Figure 10.9. VSB receiver equalizer.

The output of the forward filter and feedback filter are summed to produce the output. This output is sliced by either an 8-level slicer (15-level slicer when the comb filter is used) or a binary slicer depending upon whether the data eyes are open or not. (As pointed out in the previous Section on interference filtering, the comb filter does not close the data eyes but creates twice as many of the same magnitude). This sliced signal has the training signal and segment syncs reinserted as these are fixed patterns of the signal. The resultant signal is fed into the feedback filter, and subtracted from the output signal to produce the error estimate. The error estimate is correlated with the input signal (for the forward filter), or by the output signal (for the feedback filter). This correlation is scaled by a step size parameter, and used to adjust the value of the tap. The delay setting of the adjustable delays is controlled according to the index of the filter tap that is being adjusted.

10.2.3.8 Phase tracking loop

The phase tracking loop is an additional decision feedback loop which further tracks out phase noise which has not been removed by the IF PLL operating on the pilot. Thus, phase noise is tracked out by not just one loop, but two concatenated loops. Because the system is already frequency-locked to the pilot by the IF PLL (independent of the data), the phase tracking loop bandwidth is maximized for phase tracking by using a first order loop. Higher order loops, which are needed for frequency tracking, do not perform phase tracking as well as first order loops. Therefore, they are not used in the VSB system.

A block diagram of the phase tracking loop is shown in Figure 10.10. The output of the real equalizer operating on the I signal is first gain controlled by a multiplier and then fed into a filter which recreates an approximation of the Q signal. This is possible because of the VSB transmission method, where the I and Q components are related by a filter function which is almost a Hilbert transform. The complexity of this filter is minor because it is a finite impulse response (FIR) filter with fixed anti-symmetric coefficients and with every other coefficient equal to zero. In addition, many filter coefficients are related by powers of two, thus simplifying the hardware design.

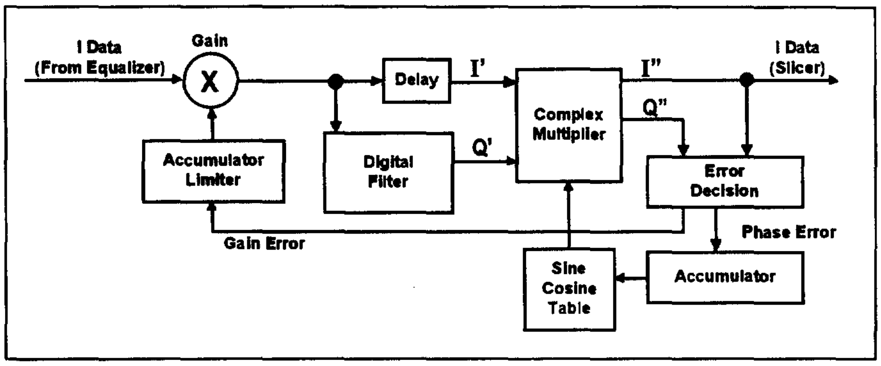

Figure 10.10. Phase tracking loop.

These I and Q signals are then fed into a de-rotator (complex multiplier), which is used to remove the phase noise. The amount of de-rotation is controlled by decision feedback of the data taken from the output of the de-rotator. As the phase tracker is operating on the 10.76 Msymbol/s data, the bandwidth of the phase tracking loop is fairly large, approximately 60 kHz. The gain multiplier is also controlled with decision feedback.

10.2.3.9 Trellis decoder

To help protect the trellis decoder against short burst interference, such as impulse noise or NTSC co-channel interference, 12 symbol code intrasegment interleaving is employed in the transmitter. As shown in Figure 10.11, the receiver uses 12 trellis decoders in parallel, where each trellis decoder sees every 12th symbol. This code interleaving has all the same burst noise benefits of a 12 symbol interleaver, but also minimizes the resulting code expansion (and hardware) when the NTSC rejection comb filter is active.

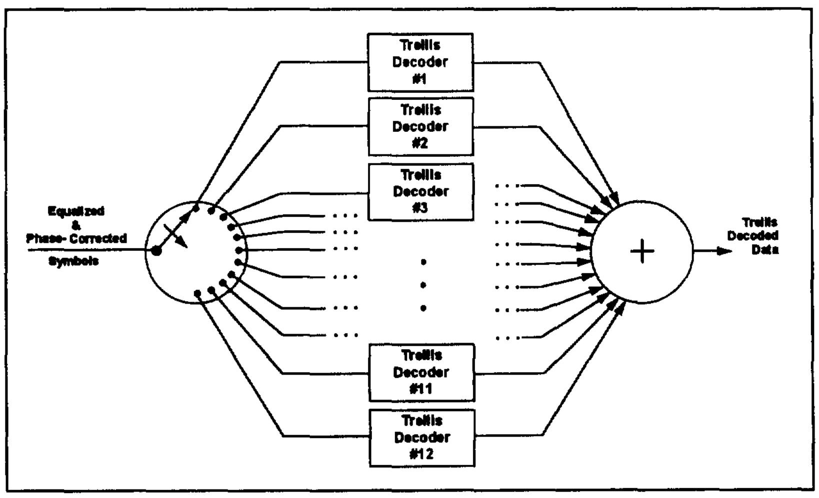

Figure 10.11. Trellis code de-interleaver.

Before the 8 VSB signal can be processed by the trellis decoder it is necessary to suspend the Segment Sync. The Segment Sync is not trellis encoded at the transmitter. The circuit block diagram which illustrates the Segment Sync suspension is shown in Figure 10.12.

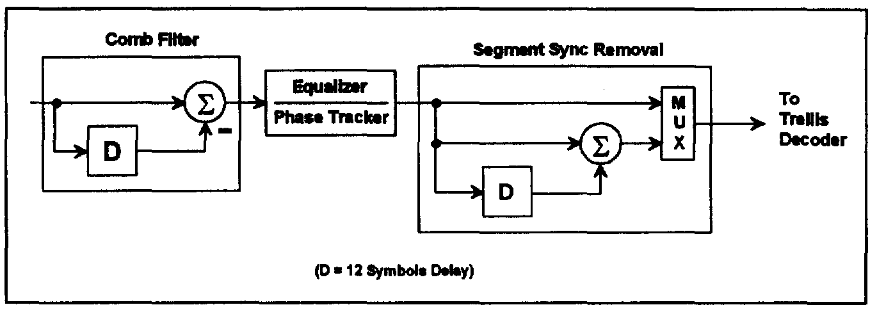

Figure 10.12. 8 VSB receiver segment sync suspension.

The trellis decoder performs the task of slicing and convolutional decoding. It has two modes; one when the NTSC rejection filter is used to minimize NTSC co-channel, and the other when it is not used. This is illustrated in Figure 10.13. The insertion of the NTSC rejection filter is determined automatically (before the equalizer), with this information passed to the trellis decoder. When there is little or no NTSC co-channel interference, the NTSC rejection filter is not used, and an optimal trellis decoder is used to decode the 4-state trellis-encoded data. Serial bits are re-created in the same order in which they were created in the encoder.

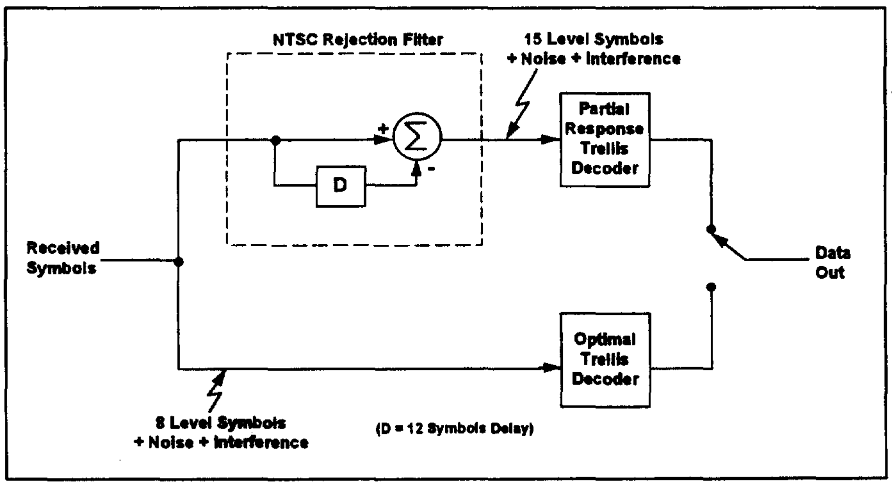

Figure 10.13. Trellis decoding with and without NTSC rejection filter.

In the presence of significant NTSC co-channel interference, when the NTSC rejection filter (12 symbol, feed-forward subtractive comb) is employed, a trellis decoder optimized for this partial response channel is used. This optimal code requires 8 states. This is necessary because the NTSC rejection filter, which has memory, represents another state machine seen at the input of the trellis decoder. In order to minimize the expansion of trellis states, two measures are taken: (1) special design of the trellis code, and (2) twelve-to-one interleaving of the trellis encoding. The interleaving, which corresponds exactly to the 12 symbol delay in the NTSC rejection filter, makes it so that each trellis decoder only sees a one-symbol delay NTSC rejection filter. By minimizing the delay stages seen by each trellis decoder, the expansion of states is also minimized. Only a 3.5 dB penalty in white noise performance is paid as the price for having good NTSC co-channel performance. The additional 0.5 dB beyond the 3 dB comb filter noise threshold degradation is due to the 12 symbol differential coding.

The presence of the Segment Sync character in the data stream passed through the comb filter presents a complication which must be dealt with because Segment Sync is not trellis encoded or precoded. Figure 10.12 shows the technique that has been used. It shows the receiver processing that is performed when the comb filter is present in the receiver. The multiplexer in the Segment Sync removal block is normally in the upper position. This presents data that has been filtered by the comb to the trellis decoder. However, because of the presence of the sync character in the data stream, the multiplexer selects its lower input during the four symbols that occur twelve symbols after the segment sync. The effect of this sync removal is to present to the trellis decoder a signal that consists of only the subtraction of two adjacent data symbols that come from the same trellis encoder, one transmitted before, and one after the segment sync. The interference introduced by the segment sync symbol is removed in this process, and the overall channel response seen by the trellis decoder is the single-delay partial response filter.

The complexity of the trellis decoder is dependent upon the number of states in the decoder trellis. Since the trellis decoder operates on an 8-state decoder trellis when the comb filter is active, this defines the amount of processing that is required of the trellis decoder. The decoder must perform an Add-Compare-Select (ACS) operation for each state of the decoder. This means that the decoder is performing 8 ACS operations per symbol time. When the comb filter is not activated, the decoder operates on a 4-state trellis. The decoder hardware can be constructed such that the same hardware that is decoding the 8-state comb filter trellis can also decode the 4-state trellis when the comb filter is disengaged, thus there is no need for separate decoders for the two modes. The 8-state trellis decoder requires less than 5000 gates.

It should be noted that after the transition period when NTSC is no longer being transmitted, the NTSC rejection filter and the 8-state trellis decoder can be eliminated from digital television receivers.

10.2.3.10 Data de-interleaver

The convolutional de-interleaver performs the exact inverse function of the transmitter convolutional interleaver. Its 1/6 data field depth, and intersegment "dispersion" properties allow noise bursts lasting about 193 μs to be handled. Even strong NTSC co-channel signals passing through the NTSC rejection filter, and creating short bursts due to NTSC vertical edges, are reliably handled due to the interleaving and RS coding process. The de-interleaver uses Data Field Sync for synchronizing to the first data byte of the data field. The convolutional de-interleaver is shown in Figure 10.14.

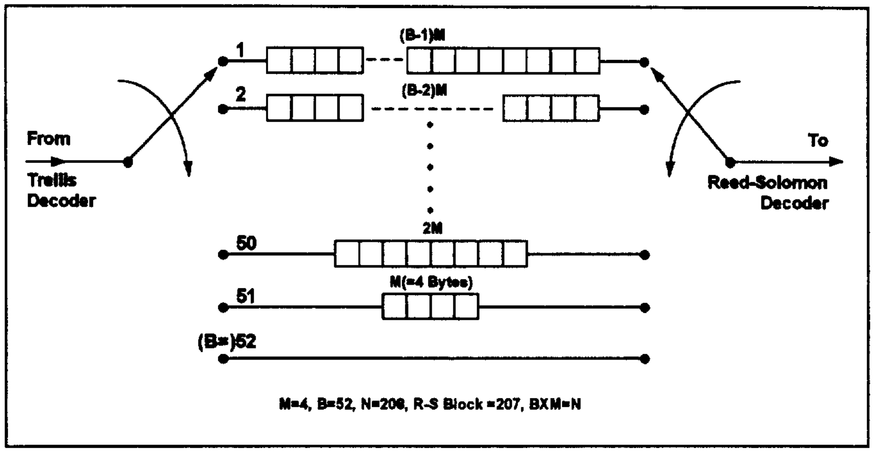

Figure 10.14. Convolutional de-interleaver.

10.2.3.11 Reed-Solomon decoder

The trellis-decoded byte data is sent to the (207,187) t=10 RS decoder, where it uses the 20 parity bytes to perform the byte-error correction on a segment-by-segment basis. Up to 10-byte errors per data segment are corrected by the RS decoder. Any burst errors created by impulse noise, NTSC co-channel interference, or trellis-decoding errors, are greatly reduced by the combination of the interleaving and RS error correction.

10.2.3.12 Data de-randomizer

The data is randomized at the transmitter by a Pseudo Random Sequence (PRS). The de-randomizer accepts the error-corrected data bytes from the RS decoder, and applies the same PRS randomizing code to the data. The PRS code is generated identically as in the transmitter, using the same PRS generator feedback and output taps. Since the PRS is locked to the reliably recovered Data Field Sync (and not some code word embedded within the potentially noisy data), it is exactly synchronized with the data, and performs reliably.

10.2.3.13 Receiver loop acquisition sequencing

The receiver incorporates a "universal reset" which initiates a number of "confidence counters" and "confidence flags" involved in the lock-up process. A universal reset occurs, for example, when tuning to another station or turning on the receiver. The various loops within the VSB receiver acquire and lock-up sequentially, with "earlier" loops being independent from "later" loops. The order of loop acquisition is as follows:

- Tuner 1st LO synthesizer acquisition
- Non-coherent AGC reduces unlocked signal to within A/D range
- Carrier acquisition (FPLL)
- Data segment sync and clock acquisition
- Coherent AGC of signal (IF and RF gains properly set)

- Data field sync acquisition
- NTSC rejection filter insertion decision made
- Equalizer completes tap adjustment algorithm
- Trellis and RS data decoding begin

Most of the loops mentioned above have confidence counters associated with them to insure proper operation. However, the build-up or let-down of confidence is not designed to be equal. The confidence counters build confidence quickly for quick acquisition times, but lose confidence slowly to maintain operation in noisy environments. The VSB receiver carrier, sync and clock circuits will work in SNR conditions of 0 dB or less as well as in severe interference situations.

10.2.3.14 High data rate cable mode

The VSB digital transmission system provides the basis for a family of ATV receivers suitable for receiving data transmissions from a variety of media. This family shares the same pilot, symbol rate, data frame structure, interleaving, Reed-Solomon coding, and synchronization pulses. The VSB system has two modes: a simulcast terrestrial broadcast mode, and a high data rate cable mode.

Most parts of the high data rate cable mode VSB system are identical or similar to the terrestrial system. A pilot, data segment sync, and data field sync are all used to provide robust operation. The pilot in the high data rate cable mode also adds 0.3 dB to the data power. The symbol, segment, and field signals and rates are all the same, allowing either receiver to lock up on the other's transmitted signal. Also, the data frame definitions are identical. The primary difference is the number of transmitted levels (8 versus 16) and the use of trellis coding and NTSC interference rejection filtering in the terrestrial system.

The high data rate cable mode receiver is identical to the VSB terrestrial receiver, except that the trellis decoder is replaced by a slicer, which translates the multi-level symbols into data. Instead of an 8-level slicer, a 16-level slicer is used. Also note that no NTSC interference rejection filter is required due to the absence of co-channel interference on cable.

10.3 Receiver equalization issues

10.3.1 Introduction

The VSB signal contains features which allow design of receivers that reliably perform the functions of acquiring and locking to the transmitted signal. The equalization of the signal for channel frequency response and ghosts is facilitated by the inclusion of specific features in the Data Field Sync (see Figure 10.15). Utilization of these features is made more reliable by the provision of means to first acquire and synchronize to the VSB signal, particularly by the Segment Sync. The Data Field Sync then can be used both to identify itself and to further perform equalization of linear transmission distortions.

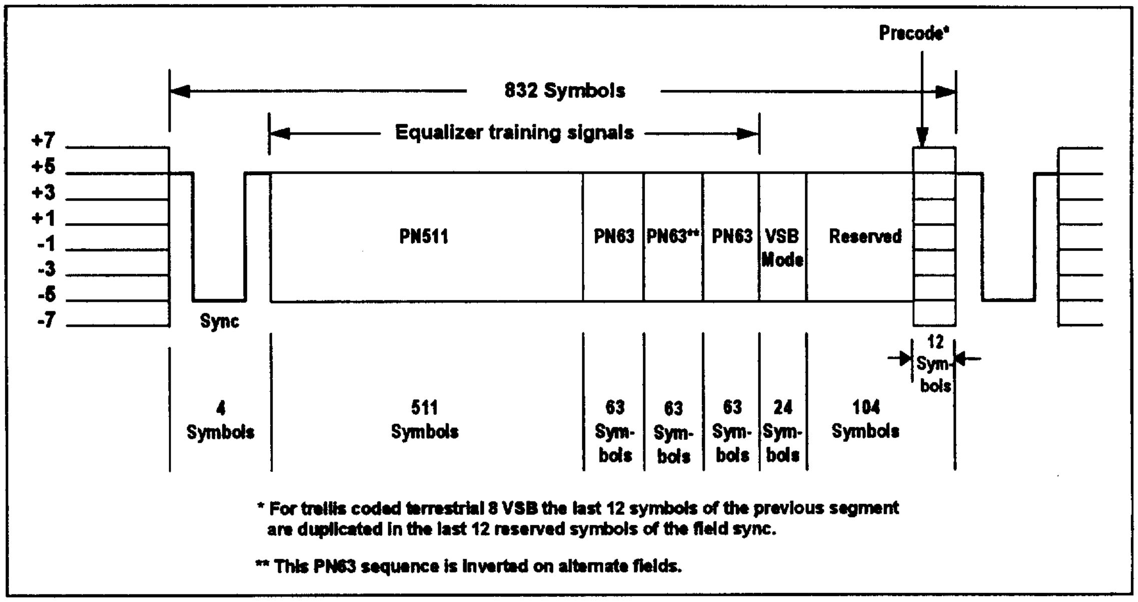

Figure 10.15. Data Field Sync.

The VSB signal may also be equalized by data-based or blind equalization methods which do not use the Data Field Sync.

10.3.2 The equalizer training signal

10.3.2.1 Specifications for the Data Field Sync

10.3.2.2 General

Data Field Sync is a unique type of Data Segment in the VSB signal. All payload data in the VSB signal is contained in Data Segments, which are processed with data interleaving, Reed-Solomon error coding, and trellis coding. The Data Field Sync (and the Segment Sync portion of every Data Segment), however, is not processed this way, as its purpose is to provide direct measurement and compensation for channel linear distortion. Equalizer training signals consisting of pseudo-noise (PN) sequences are major parts of the Data Field Sync.

10.3.3 Receiver equalization using the training signal

10.3.3.1 PN511 equalizer training sequence

This is a sequence of 511 two-level symbols (symbol rate is approximately 10.76 MHz, symbol period approximately 93 ns). The PN 511 auto-correlation function consists of one central peak surrounded by small values 30 dB or more lower than the central peak. This residual correlation noise limits the accuracy with which cancellation can be achieved to approximately 27 dB below the main signal. This is more than sufficient for the 8 VSB signal. Because the PN511 signal is not surrounded by a blank area, it will also exhibit random correlation to adjacent parts of the transmitted signal, also at a level of approximately -27 dB compared to the peak correlation. The random correlation noise can be further reduced by various means which effectively average the results of several training sequences.

10.3.3.2 PN63 equalizer training sequences

This portion of the signal consists of three identical sequential sequences of 63 two-level symbols. Because of the repetition, the correlation of one PN63 sequence with the three in the signal shows three equal peaks at 63-symbol spacing. The region of the correlation between the central peak and the outside peaks is a constant equal to -1/63 of the peak value. This allows ghosts within a delay range of 63 symbols total to be compensated exactly with one training signal measurement, as they are not affected by inherent correlation noise or intrusion of echoes from other parts of the Data Field Sync. The receiver designer may divide the total 63-symbol capability arbitrarily between pre-ghost and post-ghost correction regions. The polarity of the middle PN63 sequence is inverted on alternate fields. This removes any correlation to the PN511 sequence, if the result is averaged over an even number of fields.

10.3.3.3 Segment Sync

This is the same two-level Segment Sync used on all data segments, either payload or Data Field Sync. It is useful in equalization because it independently establishes the timing of the PN sequences and the symbol clock phase. Because of the resulting known segment timing, it is only necessary to identify which segment contains Data Field Sync. A full correlator is not required because the position of the sequence within the segment has already been established. A simple circuit which compares each segment's data to the known PN511 sequence can accomplish this identification.

10.3.3.4 VSB Mode

24 two-level symbols identify the type of data in the following data segments.

10.3.3.5 Precode

This section of twelve symbols repeats the last twelve symbols of payload data from the preceding Data Segment to allow the 12-symbol NTSC-rejection comb filter to operate properly on payload data in conjunction with trellis coding, while allowing suspension of trellis coding during the other parts of Data Field Sync.

10.3.4 Theoretical equalizer performance using training signals

10.3.4.1 Amplitude and delay of ghosts

Theoretically, in a noise-free signal, ghosts of amplitude up to 0 dB with respect to the largest signal, and within a total window of 63 symbols, can be exactly canceled in one pass using the information in the sequence. Ghosts of any delay can be canceled using the information in the sequence, with a single-pass accuracy of approximately -27 dB, which improves by averaging over multiple passes. Operation of a complete receiver system with 0 dB ghosts may not be achievable due to failure of carrier acquisition, although operation of the equalizer itself at 0 dB may be demonstrable under test by supplying an external carrier.

Practical implementations of receivers will differ in performance depending on the particular techniques used (see discussion on receiver implementation in Sections 10.3.5 and 10.3.6). The complete prototype receiver demonstrated in field and laboratory tests showed cancellation of -1 dB amplitude ghosts under otherwise low-noise conditions, and -3 dB amplitude ghost ensemble with noise. In the latter case, the signal was 2.25 dB above the receivers noise-only reception threshold.

10.3.4.2 Multiple ghosts

The number of ghosts to be canceled has in itself little effect on the theoretical performance of the system. Theoretical limits to cancellation depend on the amount of noise gain that occurs in compensating the frequency response due to the particular ghosted signal. This is affected non-linearly by the coring effect of signal quantization and decision feedback (see discussion of decision feedback in Section 10.3.6.3).

10.3.4.3 Speed of operation

Because the equalizer training signals recur with a period of approximately 24 ms, the receiver cannot perform equalization updates at a faster rate. The signal thus provides information such that the equalization system can theoretically have a bandwidth of 20 Hz. Additional constraints are implied by the desire to average out the correlation between the PN511 sequence and the alternating-polarity middle PN63 sequence.

Speed of convergence is not the only important criterion of performance, however. Ultimate accuracy and response to noise are also of importance. Equalization techniques generally proceed by successive approximation to the desired state, and therefore exhibit a convergence time measured as a number of frame periods. By using the information in the signal in different ways, speed of operation can be traded off for accuracy and noise immunity. The success of such a trade-off may depend on non-linear techniques; for example, switching between a quick acquisition mode and slower refinement mode, or using varying step sizes in a steepest-descent technique. Techniques for affecting the rate of convergence and its final accuracy are mentioned in the literature.

10.3.5 Receiver implementation using blind equalization

This section describes blind equalization techniques, which are not based on a training signal reference. They may be appropriate for use with the VSB transmission system. Blind equalization techniques are particularly useful when the channel impairments vary more rapidly than the transmission of the training wave form.

As in many modern digital data communication systems, an adaptive equalizer is used in the Grand Alliance system to compensate for changing conditions on the broadcasting channel. In communication systems which use an adaptive equalizer, it is necessary to have a method of adapting the equalizer's filter response to adequately compensate for channel distortions. There are several algorithms available for adapting the filter coefficients. The most widely used is the Least Mean Squares (LMS) method.[13]

When the equalizer is first started, the tap weights are usually not set to adequately compensate for the channel distortions. In order to force initial convergence of the equalizer coefficients, a known training signal (i.e., both the transmitter and receiver know the signal) is used as the reference signal. The error signal is formed by subtracting a locally generated copy of the training signal from the output of the adaptive equalizer. When using the training signal, the eye diagram is typically closed. The training signal serves to open the eye. After adaptation with the training signal, the eye has opened, and the equalizer may be switched to a decision directed mode of operation. The decision directed mode uses the symbol values at the output of the decision device instead of the training signal.

A problem arises in the above scenario when a training signal is not available. In this case, a method, typically called blind equalization, of acquiring initial convergence of the equalizer taps and forcing the eye open is necessary without a training signal. Blind equalization has been extensively studied for quadrature amplitude modulated (QAM) systems. For QAM systems, there are several methods typically employed: the constant modulus algorithm (CMA), and the reduced constellation algorithm (RCA) are among the most popular.[14] For VSB systems both of these methods are not directly applicable. CMA relies on the fact that, at the decision instants, the modulus of the detected data symbols should lie on one of several circles of varying diameters. Thus, it inherently relies on the underlying signal to be two-dimensional. Since VSB is essentially a one-dimensional signal (at least for the data carrying portion), CMA is not directly applicable. RCA relies on forming "super constellations" within the main constellation. The data signal is first forced to fit into a "super constellation," and then the "super constellations" are subdivided to include the entire constellation. Again, as typically used, RCA implies a two-dimensional constellation.

However, blind equalization can be performed on the VSB constellation using a modified reduced constellation algorithm (MRCA). The key part of this modification is to realize that there exists a one-dimensional version of the RCA algorithm which is appropriate for VSB signals. The MRCA consists of an algorithm to determine appropriate decision regions for a VSB decision device, so as to generate decisions which allow an adaptive equalizer to converge without the use of a training signal.

In VSB systems, the decision regions typically span one data symbol of the full constellation and the upper and lower bound of each decision region are set midway between the constellation points. If these decision regions are used for initial convergence of the equalizer, the equalizer will not converge, since, due to the presence of intersymbol interference, a significant amount of the decisions from the decision device will be incorrect .

---

[13] See [1] in the bibliography of Section 10.3.6.6.

[14] See [2] and [3] in the bibliography of Section 10.3.6.6.

In order to force more correct decisions to be made, an algorithm for determining new upper and lower decision region boundaries has been determined. The algorithm clusters the full VSB constellation into several sets, determines upper and lower bounds for decision regions, and appropriate decision device output "symbol" values. These first sets are further divided into smaller sets until each set of symbols contains exactly one symbol, and the decision regions correspond to the standard decision regions for VSB described above. The function of each stage is to allow for more decisions to be correct, and thereby drive the equalizer towards convergence. In this way, each stage in the blind equalization process serves to further open the eye.

In general, the MRCA algorithm consists of clustering the decision regions of the decision device into finer portions of the VSB constellation. The method starts with a binary (two-level) slicer, then switches to a four-level slicer, then eight-level slicer, etc. It should be noted that the MRCA algorithm is applicable to both linear equalization and decision feedback equalization.

10.3.6 Equalizer hardware implementation

10.3.6.1 General

Various system configurations need to be considered and evaluated when designing an adaptive equalizer. Some pertinent items are touched on here; references are given to further information which has appeared in the literature. A particular practical implementation will combine choices from various techniques available for achieving the desired performance, and will be based on current technology and economics as well as theoretical considerations. Choices between the alternatives presented here are not required by the Digital Television Standard, although the particular combination of techniques used in the Grand Alliance prototype system is described for reference.

10.3.6.2 Data and filter coefficient resolution

A consideration which needs brief mention is the resolution of the quantized data and the resolution of the filter coefficients. Obviously, more resolution means better performance, but at increased cost. Generally speaking, it has been found desirable to have data resolution in the range of 8 bits or more, and coefficient resolution in the range of 9 bits or more.

10.3.6.3 Equalizer filter implementation

In the design of equalizers a trade-off is usually made between Finite Impulse Response (FIR) filters vs. Infinite Impulse Response (IIR) filters.

In equalizers, FIR filters have the attractive properties of always being stable and being able to cancel both pre and post ghosts. Unfortunately, in the process of canceling a ghost they generate a secondary ghost, with twice the delay of the original, but with an amplitude that is the square of the original. This secondary ghost is then canceled in a similar manner, with a tertiary ghost being generated. This process continues until the n^{th} ghost generated is so small as to be lost in the noise of the signal, or is lost in the quantization process of the digital filtering, provided the FIR filter is long enough. To cancel long delay ghosts, FIR filters therefore need a great number of coefficients. Also, since the cancellation process involves delaying, scaling, and adding copies of the original signal back to itself, noise enhancement generally results.

IIR filters cancel ghosts without other ghosts being generated, and so they cancel long ghosts with fewer coefficients. With a slicer placed in the feedback path (decision feedback), white noise is removed from the feedback signal (assuming the slicer doesn't make any mistakes), and so the feedback signal does not increase the overall noise level of the output.

IIR filters cannot cancel pre ghosts, and, in actual applications, cannot cancel short post delay ghosts due to latencies inherent in the filtering hardware. Also, care must be taken to insure stability. A useful approach is to use an equalizer that has both an FIR and an IIR section. Here, small clumps of coefficients are positioned, with the help of programmable bulk delays, where they will be most effective in canceling the ghosts. Very long delay ghosts can then be canceled with relatively few coefficients.

10.3.6.4 Real-only vs. complex equalizer structures

In QAM systems, 2 carriers (I and Q), offset 90 degrees from one another, are modulated with independent data streams. Two equalization paths need to be employed in order to correctly equalize this signal. In the 8/16 VSB system, only one carrier (I) is modulated with unique data. The resulting Q carrier is dependent upon the I carrier, and is related to it by a mathematical function that is very similar to a Hilbert transform. (Had the 8/16 VSB system been single sideband (SSB), the Q channel would be identically equal to the Hilbert transform of the I channel. The 8/16 VSB system is not SSB, but it is very close.) Therefore, only one equalization path is needed in the 8/16 VSB system.

10.3.6.5 Symbol-rate vs. fractional-symbol sampling; use of Segment Sync

Many digital transmission systems do not use any kind of explicit synchronization information for clock recovery, and so must operate on the transmitted data itself in order to regenerate clocks. Any phase error in the sampling clock will result in degraded performance at the equalizer output. One means to overcome this is to over-sample the data (fractionally spaced sampling). A common over-sampling rate is 2 times (designated T/2), so that 2 samples are taken for every transmitted symbol. T/2 sampling reduces the effect of sampling phase error in the recovered data, but it doubles the number of filter coefficients required for a given ghost coverage range. The 8/16 VSB system periodically transmits a fixed pattern of 4 symbols (Segment Sync), and these can be used to set the phase of the sampling clock. The phase error of this clock is then extremely low, making a symbol spaced (T) equalizer very effective for equalization.

10.3.6.6 Bibliography on equalization

[1] Qureshi, Shahid U.H., "Adaptive Equalization," *Proceedings of the IEEE*, Vol. 73, No. 9, September 1985, pp. 1349-1387.

[2] Ciciora, Walter, et. al., "A Tutorial on Ghost Canceling in Television Systems," *IEEE Transactions on Consumer Electronics*, Vol. CE-25, No. 1, February, 1979, pp. 9-44.

[3] Ungerboeck, Gottfried, "Fractional Tap-Spacing Equalizer and Consequences for Clock recovery in Data Modems," *IEEE Transactions on Communications*, Vol. COM-24, No. 8, August 1976, pp. 856-864.

[4] Keeler, R. J., "Construction and Evaluation of a Decision Feedback Equalizer", *Rec. IEEE Int. Conference on. Commun.*, Philadelphia, PA, June 19-21, 1972.

[5] Widrow, B., and Stearns, S., *Adaptive Signal Processing*, Prentice Hall, 1985.

[6] Haykin, S., *Adaptive Filter Theory*, Prentice Hall, 1991.

[7] Proakis, John G., *Digital Communications*, second edition, McGraw-Hill, New York, 1989.

[8] Godard, D. N. "Self-recovering equalization and carrier tracking in two-dimensional data communication systems," *IEEE Transactions on Communications*, Vol. COM-28, No. 11, pp. 1867-1875, November, 1980.

[9] Godard, D.N., and Thirion, P.E. "Method and device for training an adaptive equalizer by means of an unknown data signal in a QAM transmission system," U.S. Patent 4 227 152, October 7, 1980.

10.4 Receiver video issues

10.4.1 Video formats

Transmissions conforming to the Digital Television Standard are expected to include the video formats as described in Table 10.3. Receivers will have to extract the picture rates and video format information, and will have to perform the necessary interpolation and format conversion so that these video formats can be displayed in the 'native' display format of the receiver.

Two native display formats were implemented by the Grand Alliance. Receivers were implemented using 787.5 scan lines per vertical scan in progressive mode and using 562.5 lines per vertical scan in interlaced mode (1125 lines per frame).

Table 10.3 The Grand Alliance Video Formats

Pixels horizontal	Pixels vertical	Pictures / sec	Interlace / progressive
1920	1080	24 (or 23.98)	Progressive
1920	1080	30 (or 29.97)	Progressive
1920	1080	60 (or 59.94)	Interlaced
1280	720	24 (or 23.98)	Progressive
1280	720	30 (or 29.97)	Progressive
1280	720	60 (or 59.94)	Progressive

10.4.2 Multiple video programs

In the case of multi-video transmissions on cable, the receiver designers should consider that the ATV display may have to be capable of also displaying video in the 525 line scanning format. It will be necessary to identify and extract the audio stream which corresponds to the user selected video stream. The use of packet ID (PID) in the header as a means of bit stream identification makes it possible to have a mix of video, audio and auxiliary data which is flexible. The control bit stream contains the program_map_table that describes the elementary stream map.

10.4.3 Concatenation of video sequences

The video coding specified in the Digital Television Standard is based on the ISO/IEC Standard 13818-2 (MPEG-2 Video). MPEG-2 Video specifies a number of video related parameters in the sequence header such as profile and level, VBV size, maximum bit rate, field/frame rate information, all progressive scan indicator, horizontal and vertical resolution, picture structure, picture aspect ratio, color field identification, chroma format, colorimetry, pan & scan etc., all of which will be necessary for receivers.

MPEG-2 Video specifies the behavior of a compliant video decoder when processing a single video sequence. A coded video sequence commences with a sequence_start_code symbol, contains one or more coded pictures and is terminated by a sequence_end_code symbol. Parameters specified in the sequence header are required to remain constant throughout the duration of the sequence. Specification of the decoding behavior in this case is feasible because the MPEG-2 Video standard places constraints on the construction and coding of individual sequences. These constraints prohibit channel buffer overflow/underflow as well as coding the same field parity for two consecutive fields.

It is envisioned that it will be common for coded bit streams to be spliced for editing, insertion of commercial advertisements, and other purposes during the video production and distribution chain. If one or more of the sequence level parameters differ between the two bit streams to be spliced, then a sequence_end_code symbol must be inserted to terminate the first bit stream, and a new sequence header must exist at the start of the second bit stream. Thus the situation of concatenated video sequences arises.

While the MPEG-2 Video standard specifies the behavior of video decoders when processing a single sequence, it does not place any requirements on the handling of concatenated sequences. Since the MPEG-2 Video standard does not specify the behavior of decoders in this case, channel buffer overflow/underflow could occur at the junction between two coded sequences unless well-constrained concatenated sequences are produced.

While it is recommended, the Digital Television Standard does not require the production of well-constrained concatenated sequences as described in Section 5.13. If well-constrained concatenated sequences are produced according to these recommendations, then it is recommended that receivers provide a seamless presentation across such concatenated sequences. Seamless presentation occurs when each coded picture is correctly decoded and displayed for the proper duration.

10.4.4 D-frames

The MPEG family of video coding standards (ISO 11172-2 and ISO 13818-2) includes a provision for efficiently coding reduced resolution pictures in "D-frames" by using intraframe DCT DC coefficients. The use of D-frames was envisioned as a means of storing highly compressed intraframe coded pictures for allowing crude fast scan of compressed video stored on digital storage media. The Digital Television Standard does not include syntax for the transmission of D-frame coded pictures; however, receivers may support the decoding of D-frames for all picture formats to allow for the use of this efficient coding mode by VCRs, digital videodisc players, or other digital storage media.

10.4.5 Adaptive video error concealment strategy

10.4.5.1 Error concealment requirements

In MPEG video compression, video frames to be coded are formatted into sequences containing intra-coded (I), predictive-coded (P) and bidirectionally predictive-coded (B) frames. This structure of MPEG implies that if an error occurs within I-frame data, it will propagate for a number of frames. Similarly, an error in a P-frame will affect the related P and B-frames, while B-frame errors will be isolated. Therefore, it is desirable to develop error concealment techniques to prevent error propagation and, consequently, to improve the quality of reconstructed pictures.

There are two approaches which have been used for I-frame error concealment, temporal replacement and spatial interpolation. Temporal replacement can provide high resolution image data as the substitute to the lost data; but in motion areas a significant difference might exist between the current intra-coded frame and the previously decoded frame. In this case, temporal replacement will produce large distortion unless some motion-based processing can be applied at the decoder. However, this type of processing is not always available since it is a computationally demanding task. In contrast, a spatial interpolation approach synthesizes the lost data from the adjacent blocks in the same frame. In spatial interpolation the intra-frame redundancy between blocks is exploited, while a potential problem of blurring remains due to insufficient high order DCT coefficients for active areas.

10.4.5.2 Error concealment implementation

To address this problem, an adaptive error concealment technique has been developed. In this scheme, temporal replacement or spatial interpolation should be used based on easily obtained measures of image activity from the neighboring macroblocks, i.e., the local motion and the local spatial detail. If local motion is smaller than spatial detail, the corrupted blocks belong to the class on which temporal replacement is applied; when local motion is greater than local spatial detail, the corrupted blocks belong to the class which will be concealed by spatial interpolation. The overall concealment procedure consists of two stages. First, temporal replacement is applied to all corrupted blocks of that class through the whole frame. After the temporal replacement stage, the remaining unconcealed damaged blocks are more likely to be surrounded by valid image blocks. A stage of spatial interpolation is then performed on them. This will now result in less blurring, or the blurring will be limited to smaller areas. Therefore, a good compromise between distortion and blurring can be obtained. This algorithm uses some simple measures, obtainable at the decoder, to adapt between spatial and temporal concealment modes. It is noted that the same idea can be used for intra-coded macroblocks of P and B-frames. The only modification is that the motion-compensation should be applied to the temporal replacement and the motion vectors (if lost) are assumed from ones in the top and bottom macroblocks.

Several new methods have been developed to further improve the accuracy of concealment. The first is a spatial concealment algorithm using directional interpolation. This algorithm utilizes spatially correlated edge information from a large local neighborhood of surrounding pixels and performs directional or multi-directional interpolation to restore the missing block.

The second method is I-picture vectors. Motion information is very useful in concealing losses in P and B pictures, but is not available for I-pictures. If motion vectors are made available for all MPEG pictures (including I-pictures) as an aid for error concealment, good error concealment performance can be obtained without the complexity of adaptive spatial processing. Therefore, a syntax extension has been adopted where motion vectors can be transmitted in an I-picture as the redundancy for error concealment purposes.

The third algorithm is the enhancement version of the adaptive spatio-temporal algorithm. The basic idea of this algorithm is to use a weighted average of spatial and temporal information rather than exclusively using either spatial or temporal information alone to conceal missing blocks. The temporal replacement estimate is further enhanced by the use of sub-macroblock refined motion vectors. Rather than applying a single estimated motion vector on an entire macroblock to create a temporal replacement (which can often result in a blocky shearing effect), every small sub-macroblock pixel region (e.g. 2x2 or 4x4 pixel regions) that composes the entire macroblock undergoes temporal replacement with its own estimated motion vector. The motion vectors associated with each sub-macroblock region are obtained from a smooth interpolation of the motion vector field, resulting in a temporal replacement estimate that is continuous at macroblock boundaries and fits well with its neighboring macroblocks.

10.4.5.3 Bibliography on error concealment

[1] H. Sun and J. Zdepski, "Adaptive Error Concealment Algorithm for MPEG Compressed Video," *Proceedings of SPIE, Visual Communication and Image Signal Processing*, Vol. 1818, Nov. 1992, pp. 814-824.

[2] H. Sun, J. Zdepski, W. Kwok and D. Raychaudhuri, "Error Concealment Algorithms for Robust Decoding of MPEG Compressed Video," Submitted to *IEEE Trans. on Circuits and Systems for Video Technology*.

[3] W. Kwok and H. Sun, "Multi-directional Interpolation for Spatial Error Concealment," *IEEE Trans. on Consumer Electronics*, Vol. 39, Number 3, August, 1993, pp. 455-460.

[4] H. Sun and W. Kwok, "Concealment of Damaged Block Transform Coded Images Using Projections onto Convex Sets," will appear in the April 1995 issue of *IEEE Trans. on Image Processing*.

10.5 Receiver audio issues

This Section summarizes receiver implementation issues related to audio. Further information on the audio system may be found in Chapter 6 of this Guide which contains information of interest to both broadcasters and receiver manufacturers.

10.5.1 Audio coding

The audio specification may be found in Annex B of the Digital Television Standard. The audio is encoded with the AC-3 system which is documented in detail in the Digital Audio Compression (AC-3) Standard, ATSC Doc. A/52. The parameters of AC-3 used in the digital television system are constrained by limiting the audio sample rate to 48 kHz, prohibiting the use of the 1+1 audio coding mode, and restricting the maximum bit rates used for audio services. A main audio service is encoded at a bit rate less than or equal to 384 kbps. A single-channel associated service is encoded at a bit rate less than or equal to 128 kbps. A two-channel associated service is encoded at a bit rate less than or equal to 192 kbps. The combined bit rate of a main service and an associated service, which are intended to be decoded simultaneously, is less than or equal to 512 kbps.

10.5.2 Audio channels and services

In general, a complete audio program which is presented to a viewer may consist of audio program elements from more than one audio elementary stream. Program elements are delivered in elementary streams tagged as to audio service type. There are eight types of audio services defined.

Two service types are defined as main audio services: complete main program (CM); and music and effects (ME). Six service types are defined as associated audio services: visually impaired (VI); hearing impaired (HI); dialogue (D); commentary (C); emergency announcement (E); and voice-over (VO). A complete audio program is constructed by decoding one complete main audio service (CM), or by decoding and combining one main audio service (CM or ME) and one associated audio service (VI, HI, D, C, E, or VO).

An audio descriptor in the PSI data provides the receiver information about the audio service types which are present in a broadcast. The transport decoder is responsible for selecting which audio service(s) elementary bit stream(s) to deliver to the audio decoder.

A main audio service may contain from one to 5.1 audio channels. The 5.1 channels are left (L), center (C), right (R), left surround (LS), right surround (RS), and low frequency enhancement (LFE). Decoding of the LFE, channel is receiver optional. The LFE channel provides non-essential low frequency effects enhancement, but at levels up to 10 dB higher than the other audio channels. Reproduction of this channel is not essential to enjoyment of the program, and can be perilous if the reproduction equipment can not handle high levels of low frequency sound energy. Typical receivers may thus only decode and provide five audio channels from the selected main audio service, not six (counting the 0.1 as one).

An associated audio service, with one exception, contains only a single audio channel. In order to simultaneously decode a main service and an associated service, it is necessary for the audio decoder to be able to decode six audio channels (five from a main service plus one from an associated service). Receivers which also support optional decoding of the LFE channel thus need to support the decoding of seven audio channels. In the case that an ME main audio service is limited to two audio channels (2/0 mode), the D service may also contain two audio channels (2/0 mode). (This exception only requires the decoding of four audio channels and thus entails no additional decoder complexity.)

It is not necessary for every receiver to completely decode all of the encoded audio channels into separate audio signals. For instance, a monophonic receiver only needs to provide a single output channel. While the single monophonic output channel must represent a mix down of all of the audio channels contained in the audio program being decoded, simplifications of the mono decoder are possible. For instance, only a single output buffer is required so that decoder memory requirements are reduced; and, some of the mixdown may occur in the frequency domain thus reducing the complexity of the synthesis filter bank.

10.5.3 Loudness normalization

There is no regulatory limit as to how loud dialogue may be in an encoded bit stream. Since the digital audio coding system can provide more than 100 dB of dynamic range, there is no reason for dialogue to be encoded anywhere near 100% as is commonly done in NTSC television. However, there is no assurance that all program channels, or all programs or program segments on a given channel, will have dialogue encoded at the same (or even similar) level. Encoded AC-3 elementary bit streams are tagged with an indication of the subjective level at which dialogue has been encoded. The receiver should be capable of using this value to adjust the reproduced level of audio programs so that different received programs have their spoken dialogue reproduced at a uniform level. The receiver may then offer the viewer an audio volume control calibrated in absolute sound pressure level. The viewer could dial up the desired SPL for dialogue, and the receiver would scale the level of every decoded audio program so that the dialogue is always reproduced at the desired level.

10.5.4 Dynamic range control

It is common practice for high quality programming to be produced with wide dynamic range audio, suitable for the highest quality audio reproduction environment. Broadcasters, serving a wide audience, typically process audio in order to reduce its dynamic range. The processed audio is more suitable for the majority of the audience which does not have an audio reproduction environment which matches that of the original audio production studio. In the case of NTSC, all viewers receive the same audio with the same dynamic range, and it is impossible for any viewers to enjoy the original wide dynamic range production.

The AC-3 audio coding system provides a solution to this problem. A dynamic range control value (dynrng) is provided in each audio block (every 5 ms). These values are used by the audio decoder in order to alter the level of the reproduced audio for each audio block. Level variations of up to +24 dB may be indicated. The values of dynrng are generated in order to provide a subjectively pleasing but restricted dynamic range. The unaffected level is dialogue level. For sounds louder than dialogue, values of dynrng will indicate gain reduction. For sounds quieter than dialogue, values of dynrng will indicate a gain increase. The broadcaster is in control of the values of dynrng, and can supply values which generated the amount of compression which the broadcaster finds appropriate.

By default, the values of dynrng will be used by the audio decoder. The receiver will thus reproduce audio with dynamic range as intended by the broadcaster. The receiver may also offer the viewer the option to scale the value of in order to reduce the effect of the dynamic range compression which was introduced by the broadcaster. In the limiting case, if the value of dynrng is scaled to zero, then the audio will be reproduced with its full original dynamic range. The optional scaling of dynrng can be done differently for values indicating gain reduction (which makes quiet sounds louder). Thus the viewer may be given independent control of the amount of compression applied to loud and quiet sounds. The details of these control functions are up to each receiver implementation.

10.6 Guide to existing TV receiver requirements and standards

10.6.1 Introduction

The following is a listing of mandatory and voluntary standards, recommended practices, and other reference information for conventional television receivers. A study of these materials will be helpful in understanding what changes will be necessary to accommodate ATV receivers and what may have to be added to cover ATV receivers, if necessary.

10.6.2 Mandatory requirements

47, CFR, (FCC) Part 2, *Marketing and Importation.*

47, CFR, (FCC) Part 15, *Radio Frequency Devices.*

FCC/OET MP-2, *Measurement of UHF Noise Figure of TV Receivers.*

21CFR (FDA), Subchapter J, *Radiological Health.*

Part 1000, *General.*

Part 1002, *Records and Reports.*

Part 1020, *Performance Standard for Ionizing Radiation Emitting Products.*

Part 1020.10, *Television Receivers.*

ANSI C63.4, *Standard Methods Of Measurement Of Radio-Noise Emissions From Low-Voltage Electrical And Electronic Equipment In The Range Of 9 kHz To 40 GHz.*

FTC-16CFR, Part 410, *Deceptive Advertising As To Sizes Of Viewable Pictures Shown By Television Receiving Sets.*

10.6.3 Mandatory for "cable ready" receivers

EIA IS-6, *Recommended Cable Television Identification Plan.*

EIA IS-23, *RF Interface Specification For Television Receiving Devices And Cable Television Systems.*

EIA-Draft IS-132, NEW, *Channelization Plan for Cable Television Tuners.*

10.6.4 Mandatory in some states of the US and/or in Canada

UL 1410, *Television Receivers And High Voltage Video Product,* (can submit receivers using this standard until 2002).

UL-1492, *Audio-Video Products And Accessories,* (can submit receivers using this standard now and beyond 2002).

UL-1413, *High Voltage Components For Television Type Appliances.*

UL-1418, *Implosion Protected Cathode Ray Tubes For Television Type Appliances.*

NFPA-70, *National Electrical Code.*

CSA-C22.2 No. 1-M90, *Radio, Television And Electronic Apparatus.*

10.6.5 Voluntary standards

EIA IS-16A, *Immunity Of Television Receivers And Video Cassette Recorders (VCRs) To Direct Radiation From Radio Transmissions, O.5 To 30 MHz.*

EIA IS-31, *Recommended Design Guideline - Rejection Of Educational Interference To Ch 6 Television Reception.*

EIA-544, *Immunity Of TV And VCR Tuners To Internally Generated Harmonic Interference From Signals.*

10.6.6 Other related standards and references

10.6.6.1 FCC

47, CFR, (FCC) Part 73, *Radio Broadcast Services.*

47, CFR, (FCC) Part 76, *Cable Television Service.*

10.6.6.2 Safety & x-ray

NOM-001, *Electronic Apparatus-Household Electronic-Apparatus By Different Sources Of Electrical Power-Safety Requirements And Testing-Methods For Type Approval.*

CPEB1, *Standard Method Of Measurement Of Ionizing Radiation From Television Receivers For Factory Quality Assurance.*

CPEB2, *Definition Of Normal Operating Conditions For Television Receivers.*

CPEB3, *Measurement Instrumentation For X-Radiation From Television Receivers.*

EIA-500-A, *Recommended Practice For Measurement Of X-Radiation From Projection Cathode Ray Tubes.*

EIA-503-A, *Recommended Practice For The Measurement Of X-Radiation From Direct View Television Picture Tubes.*

10.6.6.3 Interference & immunity

EIA-378, *Measurement Of Spurious Radiation From FM And TV Broadcast Receivers In The Frequency Range Of 100 To 1000 MHz - Using The EIA Laurel Broadband Antenna.*

CISPR 13, *Limits And Methods Of Measurement Of Radio Interference Characteristics Of Sound And Television Broadcast Receivers And Associated Equipment.*

CISPR 20, *Limits And Methods Of Measurement Of Immunity Characteristics Of Sound And Television Broadcast Receivers And Associated Equipment.*

IEEE-187-90, *IEEE Standard On Radio Receivers: Open Field Method Of Measurement Of Spurious Radiation From FM And Television Broadcast Receivers.*

10.6.6.4 Cathode ray TV display tubes

EIA -256-A, *Deflecting Yokes For Cathode Ray Tubes.*

EIA-266-A, *Registered Screen Dimensions For Monochrome Picture Tubes.*

EIA-324-A, *Registered Screen Dimensions For Color Picture Tubes.*

EIA-493, *Recommended Practice For Conversion Of U.S. To Metric Dimensions For Color And Monochrome Cathode Ray Tubes And Their Component Parts.*

EIA-527, *Screen Definition For Color Picture Tubes.*

10.6.6.5 Voluntary TV receiver recommended practices

EIA-462, *Electrical Performance Standards For Television Broadcast Demodulators.*

EIA-563, *Standard Baseband (Audio/Video) Interface Between NTSC Television Receiving Devices And Peripheral Devices.*

REC-109-CH, *Intermediate Frequencies For Entertainment Receivers.*

TVSB1, *EIA Recommended Practice For Use Of A Vertical Interval Reference (VIR) Signal.*

TVSB3, *A History Of The Vertical Interval Color Reference Signal (VIR).*

TVSB5, *Multichannel TV Sound System - BTSC System Recommended Practices.*

10.6.6.6 International standards, IEC, etc.

IEC 65, *Safety Requirements For Mains Operated Electronic And Related Apparatus For Household And Other Similar General Use.*

IEC 107-1 to -6, *Methods Of Measurement On Receivers For Television Broadcast Transmissions,* (6 Parts).

IEC 569, *Informative Guide For Subjective Tests On Television Receivers.*

IEC 68, *Environmental Testing - Part 1 : General and Guidance.*

Blank Page